arXiv:1612.04869v1 [cs.CV] 14 Dec 2016

Border-Peeling Clustering

Nadav Bar, Hadar Averbuch-Elor, Daniel Cohen-Or
School of Computer Science, Tel Aviv University
{nadavbar, averbuch1, dcor}@mail.tau.ac.il

Abstract

In this paper, we present a novel non-parametric clustering technique, which is based on an iterative algorithm that peels off layers of points around the clusters. Our technique is based on the notion that each latent cluster is comprised of layers that surround its core, where the external layers, or border points, implicitly separate the clusters. Analyzing the K-nearest neighbors of the points makes it possible to identify the border points and associate them with points of inner layers. Our clustering algorithm iteratively identifies border points, peels them, and separates the latent clusters. We show that the peeling process adapts to the local density and successfully separates adjacent clusters. A notable quality of the Border-Peeling algorithm is that it does not require any parameter tuning in order to outperform state-of-the-art finely-tuned non-parametric clustering methods, including Mean-Shift and DBSCAN. We further assess our technique on high-dimensional datasets that vary in size and characteristics. In particular, we analyze the space of deep features that were trained by a convolutional neural network.

1 Introduction

Clustering is the task of categorizing data points into groups, or clusters, with each cluster representing a different characteristic or similarity between the data points. Clustering is a fundamental data analysis tool, and as such has abundant applications in different fields of science and is especially essential in an unsupervised learning scenario. Ideally, a clustering method should infer the structure of the data, e.g., the number of clusters, without any manual supervision.

Many of the state-of-the-art clustering methods operate with several underlying assumptions regarding the structure of the data. A prominent assumption is that the clusters have a single area that can be identified as the center, or the core of the cluster. For instance, K-Means [12] operates under the assumption that there is a single cluster center according to the compactness of the data, while the Mean-Shift [2] method defines this area as the one displaying the highest density inside the cluster. Operating under this assumption may result in overly split clusters containing several dense areas, or centers, of smaller clusters.

Density based methods like DBSCAN [3] operate often under the assumptions that different clusters have similar levels of density, and that the clusters are easily separable from one another. However, often, clusters do not have such structural density and are not clearly separable, leading to a redundant merge of adjacent clusters.

In this work, we introduce *Border-Peeling*: a non-parametric clustering method that iteratively peels off layers of points until the remaining data points form areas of dense, highly separable clusters. Our technique is non-parametric in the sense that the number of clusters is not provided as input. The careful repeated peeling forms a transitive association between the peeled *border* points and the remaining *core* points. The key is to consider a layered structure for the latent clusters, where the external layers implicitly separate the clusters. We analyze the K-nearest neighbors of each point to iteratively estimate the border points and associate them with inner-layer points. See Figure 1 for an illustration our iterative technique.

In the following sections, we describe the details of our method, and demonstrate the performance of our method on various datasets and in comparison to other well-known clustering methods. Particularly, we evaluate our technique on high-dimensional datasets of deep features that were trained by convolutional neural networks. We show that the Border-Peeling clustering algorithm does not require any parameter tuning to outperform state-of-the-art non-parametric clustering methods, including Mean-Shift and DBSCAN, when the free-parameters of these are set according to their best performance.

2 The Algorithm

Given a set of points in $\mathbb{R}^d$, our clustering technique iteratively *peels off* border points from the set. The premise of this method is that by repeatedly peeling off border points, the final remaining points, termed *core points* will be better separated and easy to cluster.

To cluster the input points, during the iterative peeling

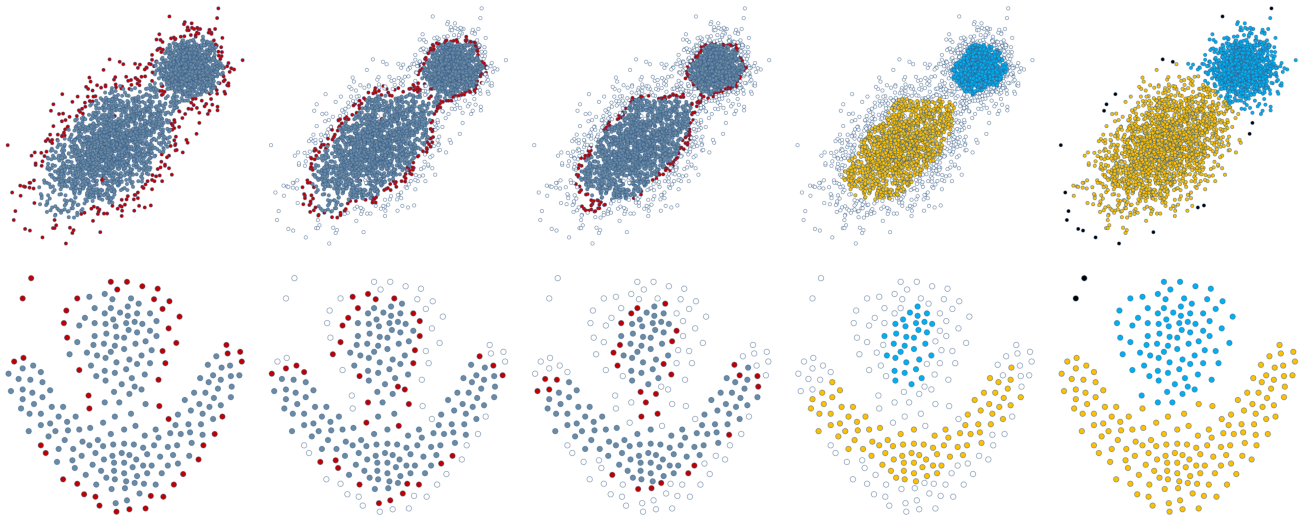

Figure 1: The Border-Peeling technique on (top row) a dataset generated from two-dimensional multivariate Gaussian distributions, and on (bottom row) a known dataset from the clustering literature [5]. Three consecutive peeling iterations are illustrated on the left, with the identified border points colored in red. Next, we illustrate the clustering of the highly-separable core points that remain after the peeling process. The rightmost figures illustrate our final clustering results.

process, each peeled point is associated and linked with a neighboring point that was not identified as a border point. This linkage forms a transitive association between the peeled points with one of the core points. The clustering of the peeled border points is then inferred by their association to the clustered core points.

The algorithmic key idea of our technique is twofold: (i) the definition of a border point, and (ii) the association of a border point to its neighboring non-border point. These key ideas will be elaborated in the following section.

The iterative peeling terminates when the identified border points are strictly weaker in terms of their "borderness" than the border points that were identified in the previous iterations, thus forming the set of core points. These core points are then grouped into clusters using a simplified version of DBSCAN.

In what follows, we first introduce some notations and describe how border points are identified at each iteration (Section 2.1). We then detail the border point association process (Section 2.2). Finally, we describe our clustering procedure (Section 2.4).

2.1 Border Points Identification

Given a set of n data points $X = \{x_1, x_2, ..., x_n\}$ in $\mathbb{R}^d$ and a similarity function $d : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ as input, we denote $X^{(t)}$ as the set of points which remain *unpeeled* by the start of the t^{th} iteration, with $X^{(1)} = X$, and $X^{(T+1)}$ representing the final set of core points for a peeling process with T iterations.

For each t, we associate the following terms with each point in $X^{(t)}$: a neighborhood $N_k^{(t)}(x_i)$, which contains the set of k nearest neighbors, a *density influence* value $b_i^{(t)}$, and a border classification value $B_i^{(t)}$ that accepts the value of 1 if x_i is a border point and 0 otherwise. Following [9] and [20], we denote the reverse k-nearest neighbors of x_i at iteration t by $RNN_k^{(t)}(x_i) = \left\{x_j | x_i \in N_k^{(t)}(x_j)\right\}$. That is, $RNN_k^{(t)}(x_i)$ is the set of points for which x_i is one of their k-nearest neighbors at the t^{th} iteration.

The *density influence* measure $b_i^{(t)}$ is defined as follows:

$$b_i^{(t)} = \sum_{x_j \in RNN_k^{(t)}(x_i)} f(x_i, x_j), \tag{1}$$

where $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ quantifies the pairwise relationship between x_i and x_j. In case that $RNN_k^{(t)}(x_i) = \emptyset$ we set $b_i^{(t)} = 0$.

In our work, for the case in which the similarity measure d is the Euclidian distance function, we calculate f by applying a Gaussian kernel with local scaling [21] to the Euclidean distance, such that:

$$f(x_i, x_j) = exp\left(-\frac{||x_i - x_j||_2^2}{\sigma_j^2}\right), \tag{2}$$

In our experiments, we followed the choice of σ_j in [21] and set $\sigma_j = ||x_j - N_k^{(t)}(x_j)[k]||_2$ where $N_k(x_j)[k]$ denotes the k-th nearest neighbor of x_j at iteration t. According to

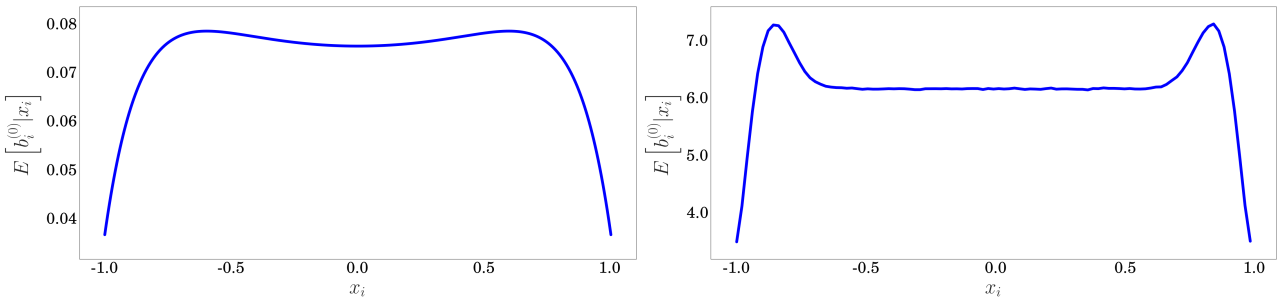

Figure 2: (A) The values of $E\left[b_i^{(0)}|x_i\right]$ for different values of x_i for $n = 50$. As expected, the minimal values are obtained at the endpoints of the range. (B) Empirical evaluation of $E\left[b_i^{(0)}|x_i\right]$ using random sampling with $n = 100$ and $k = 10$. The values of $E\left[b_i^{(0)}|x_i\right]$ were calculated by repeating the experiments and averaging over the values of $b_i^{(0)}$.

the local scaling approach, the distance to the k^{th} neighbor of the data point is used as the normalizing factor for the Gaussian kernel. This approach has proven to be effective in measuring the affinity between data points when the affinity of the data points has a large variance.

The choice for the estimation of b_i is reasonable in the sense that for each data point, it captures the amount of influence a point has on the local density of its neighboring points. Our approach can cope with multiple distribution models, as $k - NN$ neighborhoods, unlike ϵ-neighborhoods, are not affected by varying distributions in the data.

In general, we would expect the values of b_i to be smaller for data points that lie on the border of the cluster, and larger for points that are closer to the core of the clusters. Formally, we consider the following simple case.

Lemma 1. Consider a single cluster of n points in $\mathbb{R}^1$ distributed uniformly in the interval $[-1, 1]$ and let $k = 1$. Then the expected value of $b_i^{(0)}$ is given by the following expression:

$$E\left[b_i^{(0)}|x_i\right] = \left(\frac{2}{n} + \frac{n\left(x_i+1\right)-2}{2n}e^{-\frac{n}{2}(x_i+1)} + \frac{n\left(1-x_i\right)-2}{2n}e^{-\frac{n}{2}(1-x_i)}\right)\cdot e^{-1}. \quad (3)$$

A full proof of Lemma 1, together with a more detailed analysis, are available in the Supplementary Materials. The expected values described in Equation 3 in the paper are depicted in Figure 2.A. As can be seen in the figure, points close to one of the two ends of the interval (near -1 and 1) obtain, as anticipated, the smallest expected b_i values, and these are indeed the points which our algorithm would consider as border points. Figure 2.B illustrates the empirical behavior of $b_i^{(0)}$ for a more complicated scenario ($n = 100, k = 10$). As the figure illustrates, the values of $b_i^{(0)}$ continue to drop near the two ends of the interval for larger values of k.

Our border classification value $B_i^{(t)}$ is space-variant. Simply put, we learn the local characteristics of the dataset to determine whether a point is classified as a border point or not. The calculation of $B_i^{(t)}$ in each iteration is performed using an iteration specific cut-off value, denoted as $\tau^{(t)}$ such that $B_i^{(t)} = 1$ if $b_i^{(t)} \leq \tau^{(t)}$ and $B_i^{(t)} = 0$ otherwise. The set of cut-off values $\tau^{(1)}, \tau^{(2)}...$ can be manually specified, or as we describe below, it can be estimated from the intrinsic properties of the input data.

We denote the set of identified border points at iteration t as $X_B^{(t)} = \left\{x_i \colon B_i^{(t)} = 1 \wedge x_i \in X^{(t)}\right\}$, respectively, we denote the set of non-border point at iteration t as $X_{B^c}^{(t)} = X^{(t)} \setminus \cup_{r=1}^{t} X_B^{(r)}$.

2.2 Border Points Association

Following the identification of border points at iteration t, we associate to each identified border point $x_i \in X_B^{(t)}$, a neighboring non-border point which we denote as $\rho_i \in X_{B^c}^{(t)}$.

In order to prevent a scenario in which an isolated border point is associated with a relatively distant neighbor, resulting with erroneous merging of distant clusters, we use a time and spatially-variant threshold value denoted by $l^{(t)}(x_i)$. We choose ρ_i such that $\rho = argmin_{x_j \in X_{B^c}^{(t)}}\left\{d(x_i, x_j) \colon d(x_i, x_j) \leq l^{(t)}(x_i)\right\}$. That is, ρ_i is the nearest non-border point to x_i at iteration t from the set of non-border points that are in distance of less than $l^{(t)}(x_i)$ from x_i. In cases where $\left\{d(x_i, x_j) : d(x_i, x_j) \leq l^{(t)}(x_i)\right\} = \emptyset$ the value of ρ_i is left unassigned.

For each data point x_i, we calculate the value of $l^{(t)}(x_i)$ according to properties of its local neighborhood. We have found this method to perform better than when using a constant threshold value, as it takes into account the spatially-

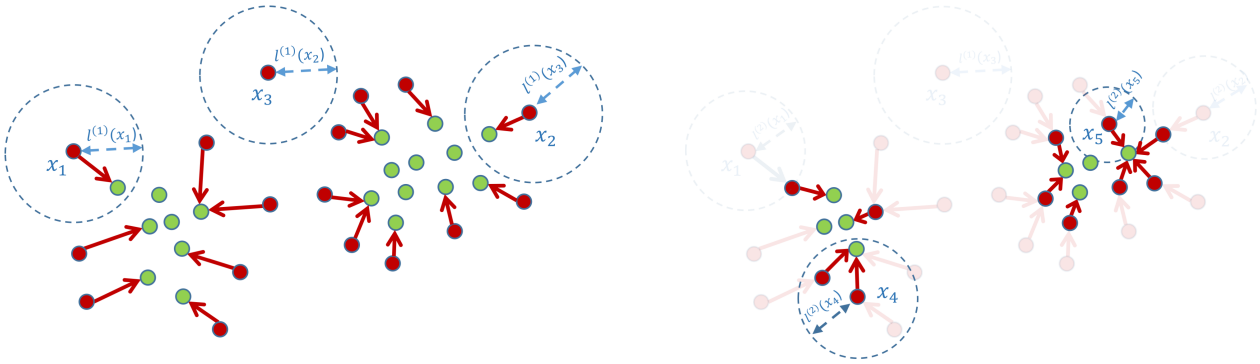

Iteration 1 Iteration 2

Figure 3: The border association process and calculation of $l^{(t)}(x_i)$ during the peeling process. The figures illustrate the border points (shown in red) of the first two iterations and their association to non-border points, marked by the red arrows. The blue circles represent the association area induced by the time and spatially-variant threshold value $l^{(t)}(x_i)$.

varying density of the data points. See Figure 3.

In the first iteration, we set $l^{(1)}(x_i) = \lambda$ for each i, where λ is a parameter of our method which serves as the maximal threshold value. In the following iterations, the value of $l^{(t)}(x_i)$ is updated in the following manner:

1. Following each iteration t in which x_i remains unpeeled, we take the k-nearest data points to x_i from the data points that are in $\cup_{r=1}^{t} X_B(r)$, i.e., the set of points that were already peeled up to the current iteration t. We denote this set by $NN_{B,k}^{(t)}(x_i)$. We then compute $l'^{(t)}(x_i)$ according to: $l'^{(t)}(x_i) = C \cdot \frac{1}{k} \sum_{x_j \in NN_{B,k}^{(t)}} l^{(t)}(x_j)$. Where the constant C determines the strictness of the threshold values. In the case that $l'^{(t)}(x_i) < \lambda$, we set $l^{(t)}(x_i) = l'^{(t)}(x_i)$.

2. Following the iteration in which x_i was peeled, if ρ_i was assigned we set $l^{(t)}(x_i) = d(x_i, \rho_i)$.

Figure 3 illustrates the effects of the time and spatially-variant threshold value $l^{(t)}(x_i)$. As the figure demonstrates, at first the association areas (illustrated by blue circles) are of equal size since at the start of the first iteration the values of $l^{(1)}$ are identical. Some data points (such as x_3 in the figure) are not associated to a non-border point since there are no non-border points with distance of at most λ. The right figure illustrates the updated value of $l^{(2)}$ for the peeled points, as the values are updated to hold the distance of each border point from its associated non-border point. The values of $l^{(2)}$ for the newly identified border points are calculated by averaging over the values of the previously peeled border points. Note, for example, that $l^{(2)}(x_4)$ is larger than $l^{(2)}(x_5)$ since the values of $l^{(2)}$ for the border points in the proximity of x_4 are larger than the values of $l^{(2)}$ for the border points that are in the proximity of x_5.

2.3 Border Peeling Termination Criterion

To automatically set the number of iterations of the peeling process T, we make the following observation: peeled border points should reside in denser areas than the border points of the previous iteration. Therefore, the assigned density influence values should increase in average (which implies that the points establish less coherent borders than the previous iterations). Furthermore, when over-clustering occurs, the values of the border points are expected be significantly higher than the values of the border points in the preceding iterations.

Hence, in each iteration t, we track the set of values of border points that are about to be peeled: $\left\{ b_i^{(t)} \middle| B_i^{(t)} = 1 \right\}$, and calculate the mean value of that set, denoted by $\bar{b}_p^{(t)}$. We then examine at the 2D graph induced by the values of $\bar{b}_p^{(1)}, \bar{b}_p^{(2)}, ..., \bar{b}_p^{(t)}$, and stop the peeling process in the case in which the ratio $\frac{\bar{b}_p^{(t)}}{\bar{b}_p^{(t-1)}}$ is significantly larger than the preceding ratios: $\frac{\bar{b}_p^{(i)}}{\bar{b}_p^{(i-1)}}$ for $i < t$. See Figure 4 for an illustration of the method.

2.4 Bottom-up Clustering

Our method iteratively identifies border points and associates them to non-border points. The transitive association thus yields paths from each input point to one of the core points, which are the final remaining non-border points, i.e., the set $X^{(T+1)}$ with the set that contains ρ_i.

When the border peeling process terminates, the remaining set of core points are clustered by merging together close reachable neighborhoods of points. Formally, a pair of core points $x_i, x_j \in X^{(T+1)}$ is said to be reachable w.l.o.g. if there is a series of core points: $(x_{k_1}, ...x_{k_m})$ with

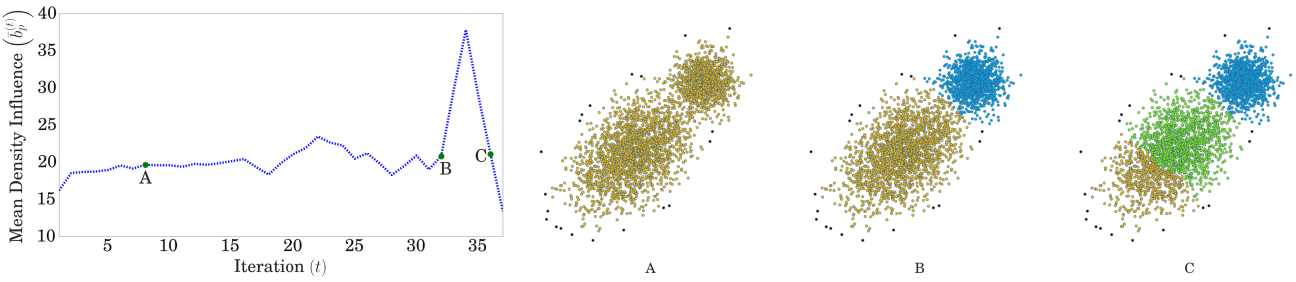

Figure 4: Automatically terminating the peeling iterations. Above, we demonstrate the 2D graph induced by the different values of $\bar{b}_p^{(t)}$, together with the corresponding clustering results, which are computed after 13 (A), 30 (B), and 36 (C) iterations. As (A) demonstrates, an early termination results in an under-clustering of the data, while too many peeling iterations (C) result in an over-clustering. The correct clustering, illustrated in (B), is achieved just before the ratio between the consecutive mean density influence values significantly increases.

$k_1 = i$ and $k_m = j$, such that for every two adjacent indices in the series (k_r, k_{r+1}) the relation $d(x_{k_r}, x_{k_{r+1}}) \leq max(l^{(T)}(x_{k_r}), l^{(T)}(x_{k_{r+1}}))$ holds. For every pair of core points (x_i, x_j) that are reachable, we merge the set that contains x_i with the set that contains x_j. This merging step is done iteratively, until all sets of reachable data points are merged.

The set of cluster candidates are then defined by following the border points association and linkage to the core points. In order to be able and better filter out noise, we mark small clusters as noise, using a user defined value of the minimum cluster size. Following the filtering step, the final set of clusters is returned.

3 Experiments

To evaluate the performance of the Border-Peeling clustering method, we measured its performance on numerous synthetic and real-life datasets and compared its performance to other state-of-the-art algorithms. The Border-Peeling clustering method was implemented in the Python programming language using the numpy [15] software library. In order to compare our performance with other clustering algorithms we used the implementation available in the SciKit-Learn [14] Python library. The time complexity of our technique is $O(T \cdot n \cdot \tilde{f}_{knn})$, where $\tilde{f}_{knn}$ is the asymptotic complexity of the $k-NN$ method.

The implementation of our method together with the datasets used in the following sections, are available at: https://github.com/nadavbar/BorderPeelingClustering.

Throughout the evaluation of Border-Peeling clustering in the experiments described below, the method was run in the following way: we fixed the Border-Peeling method parameters to $C = 3$ and set the minimum cluster size to 10 for datasets with less than 1000 data points, and 30 for larger datasets. In each iteration, $\tau^{(t)}$ was set such that 90% of the remaining data points at each iteration have larger values of $b_i^{(t)}$, and a value of $k = 20$ was used for the k-NN queries. The value of λ was calculated by first calculating all of the pairwise distances in the k-neighborhood of each point: $D_k = \cup_{x_i \in X} \{d(x_i, x_j) | x_j \in NN_k(x_i)\}$, and then setting $\lambda = MEAN(D_k) + STD(D_k)$. While other choices of λ can be used, we have found this simple estimation method to be effective, as can be seen in the results below.

3.1 Evaluation on common synthetic datasets

First, we evaluated our method on a number of common synthetic datasets from the known clustering literature [5, 6, 17], which consist of a small number of clusters that lie in proximity to one another and are not easily separable, thus constituting a challenge to density-based methods. We compare our performance to that of a set of well-known non-parametric clustering techniques: DBSCAN (DB), Mean-Shift (MS), Affinity Propagation (AP) [4] and the more recent Hierarchical-DBSCAN (HDB) [1]. Similar to Border-Peeling clustering, all of those clustering methods try to infer the number of clusters automatically and do not accept the number of clusters as a parameter. We have also compared our performance to the K-Means (KM) and Spectral Clustering (SC) [13] methods, which accept the number of clusters as a parameter. For each dataset we ran these parametric techniques 1000 times with random initialization, taking the clustering which minimizes the sum of distances of data points to their closest cluster center among those 1000 runs. We provide a comparative visualization of the clusters formed by Border-Peeling clustering to the clusters formed by K-Means (Figure 5) and Spectral Clustering (Supplementary Materials), which both obtained similar results.

We evaluated the performance of each method by cal-

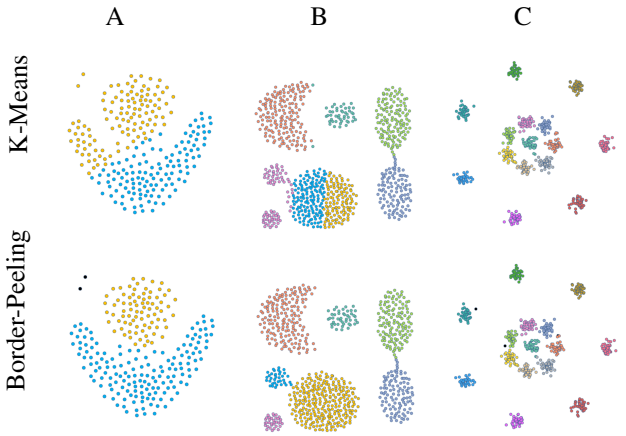

Figure 5: Qualitative comparison of our Border-Peeling technique to the parametric K-Means technique. Illustrated above are the parametric K-means (top row), where K is given as input, vs. our non-parametric clustering results (bottom) on three known synthetic examples from the literature (labeled A-C). As the figure demonstrates, Border-Peeling clustering successfully recovered the number of clusters automatically as well as identified the outliers in the synthetic datasets (colored in black), while the K-Means method failed to identify the correct clusters for data sets A and B.

culating the Adjusted Rand Index (ARI) [7], and Adjusted Mutual Information (AMI) [18] of the resulting metrics. The ARI and AMI are well-known metrics in the field of data clustering and are frequently used in order to evaluate clustering performance when ground truth data is available. The scores for the well-known clustering methods were obtained by running them over a range of parameters and taking the clustering result with the best AMI and ARI scores, while Border-Peeling clustering was run using the parameters described in the previous section.

As Figure 6 demonstrates, our technique successfully identified the number of clusters for each of the datasets and achieved the highest scores for the first two datasets, and the second highest score for the third dataset.

3.2 Evaluation on large datasets

We further evaluated the performance of the Border-Peeling method in comparison to other well-known clustering algorithms by running it on large datasets. We generated large sets by extracting feature vectors generated by convolutional neural networks (CNN) that were trained separately on MNIST [11], a handwritten digit image dataset, and on the CIFAR-10 [10] image dataset.

The MNIST dataset consists of 70000 labeled images of handwritten digits divided into a training set of 60,000 im-

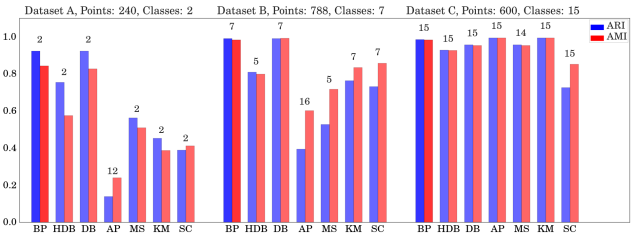

Figure 6: Comparison of Border-Peeling clustering (BP) with non-parametric (HDB, DB, AP, MS) and parametric (KM, SC) clustering techniques on the common synthetic datasets A, B and C that are illustrated in Figure 5. The number above each bar represents the number of detected clusters. Our technique (BP) detected the correct number of clusters for all three datasets, and obtained the highest score for the first two datasets and the second highest for the third dataset. The results for our method were obtained by using fixed set of parameters, while the results for the other methods were obtained by taking the best results over a large range of different parameters for the non-parametric approaches, and by taking the best result over 1000 runs for KM and SC.

ages and a test set of 10,000 images. Each image in the database is assigned to one of the possible 10 classes of digits. The CIFAR-10 dataset consists of 60000 color 32x32 images separated into 10 classes of different animals and vehicles. The dataset is divided into a training set of 50000 images and a test set of 10000 images.

In the setting of our experiment, we trained a CNN for each of the datasets using the training set and then used it to produce an embedding of the images in the test set to n-dimensional feature vectors, where $n = 500$ for MNIST and $n = 64$ for CIFAR-10. To obtain the feature vectors, we used the CNN implementations for MNIST and CIFAR-10 which are available in the MatConvNet [16] library. To produce less balanced datasets with an unknown number of clusters, the embedded test sets were sampled by taking all the embedded vectors which are within a certain radius of randomly sampled points. By using different radii, we generated several datasets of varying sizes. Embeddings of selected samples of the datasets to 2D can be seen in the top part of Figure 7. As a final preprocessing step, we employed dimensionality reduction on the sampled datasets using PCA [8] to 30 dimensions for the MNIST samples, and spectral embedding [19] to 8 dimensions for the CIFAR-10 samples.

The results of running Border-Peeling clustering as well as the other non-parametric clustering techniques on the MNIST and CIFAR-10 processed samples are shown in the bottom part of Figure 7. We performed ten different runs for each of the radii values that are used for the sampling

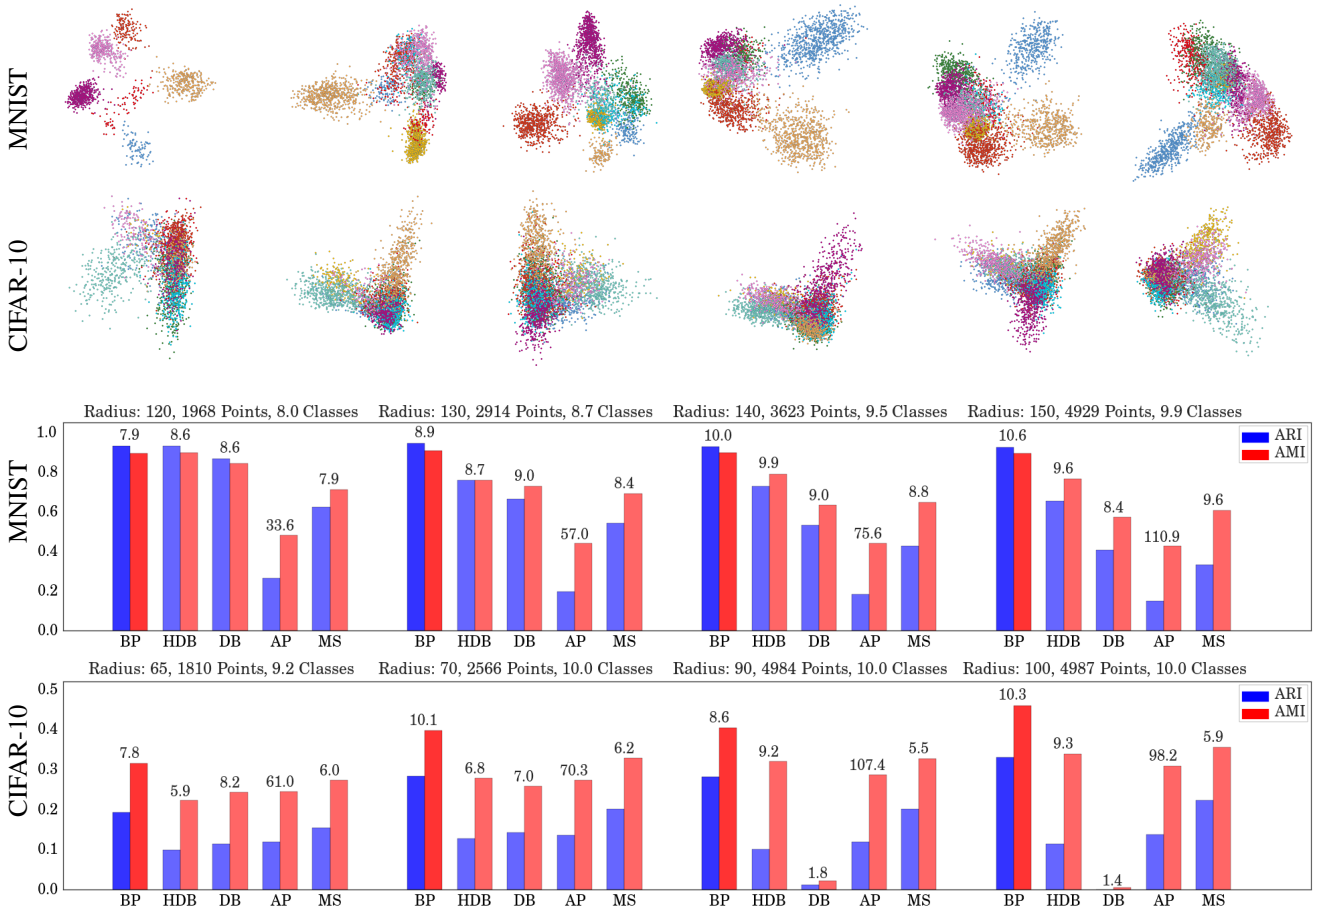

Figure 7: The top part of the figure illustrates 2D embeddings for selected samples of the MNIST and CIFAR-10 features that were produced by a CNN. The embeddings were obtained by running PCA on the datasets. As the figure illustrates, the data has no clear clusters that are easy to separate. The graphs at the bottom of the figure show a comparison with non-parametric clustering techniques on samples of the MNIST and CIFAR-10 datasets. The number above each bar represents the number of detected clusters. For both datasets our technique (BP) achieved the highest ARI and AMI scores on average.

and average the results of the clustering methods over these random runs. Throughout our experiments, the Border-Peeling parameters were set as described above, while the parameters of the other techniques were set separately for the MNIST and CIFAR-10 datasets by taking the parameters that yield the best overall result in average. As Figure 7 demonstrates, Border-Peeling clustering achieved the best performance on average in terms of AMI and ARI for both datasets. For the samples belonging to the MNIST dataset, the performance of the clustering is also comparable in the global sense, as it maintained an overall score of $ARI > 0.9$ on average for the different samples. The CIFAR dataset is more challenging, and the overall ARI and AMI scores do not surpass 0.5. However, in comparison to the other clustering methods, our technique achieves the best scores in all settings. Additionally, it is also the only method among the different clustering methods that is able to maintain good performance on a larger sampling of the datasets.

Our incremental peeling process can infer a confidence value associated with the data points of each cluster. As discussed in Section 2.1, data points with lower values of b_i are expected to be along the border of the cluster, and thus, with a lower confidence. To illustrate this simple confidence ranking, Figure 8 visualizes the top-10 and bottom-10 instances of clustering the MNIST dataset. Note that the images of digits with lower values of $b_i^{(0)}$ are often harder to identify, while images with higher values are clearer and more pronounced.

3.3 Sensitivity analysis

To demonstrate that our method is insensitive to small variations of its automatically estimated parameters, we

Figure 8: The images corresponding to the bottom-10 (odd rows) and top-10 values (even rows) of $b_i^{(0)}$ for four different clusters which were obtained using Border-Peeling clustering on a random subset of the MNIST dataset.

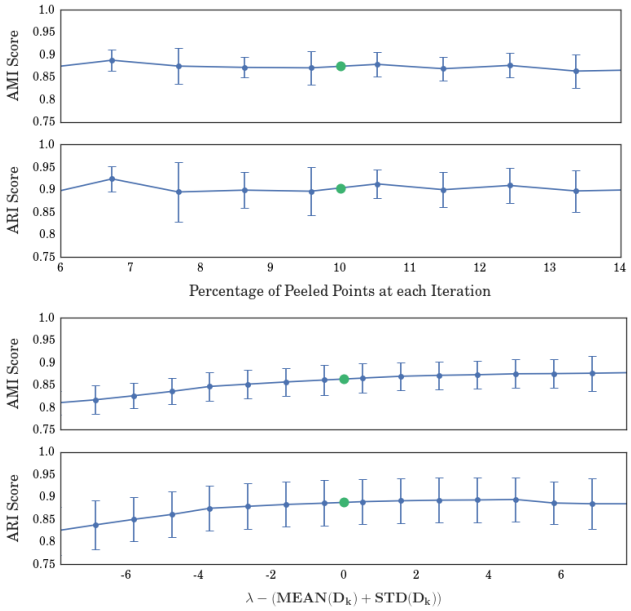

Figure 9: The ARI and AMI scores of the different clustering results on the MNIST dataset that were obtained using multiple runs of varying values of our specified parameters. The average scores are illustrated with blue circles together with the corresponding standard deviation. The parameter values that were used in the experiments section are marked with a green circle. As the figure demonstrates, both scores remain stable as the parameters sway from the default values.

performed multiple runs on the MNIST CNN features over a wide range of different parameter values. We first produced 10 different samples of the MNIST features using the method described above with a radius of 150. Each sample was then clustered using different values of λ, in the range of $[(MEAN(D_k) + STD(D_k)) - 7, (MEAN(D_k) + STD(D_k)) + 7]$. Additionally, we tested the effect of the percentage of points that are peeled at each iteration by clustering using different percentile values when choosing $\tau^{(t)}$, with values ranging between 6% and 14%. We then computed the AMI and ARI scores of the clustering results and averaged the scores for the each of the different parameter values. The results of the different runs are illustrated in Figure 9. As can be seen in the figure, both the ARI and AMI scores remain stable across the different parameter values, which affirms that our method is insensitive to its estimated parameter values.

4 Conclusions

We have presented a non-parametric clustering technique that groups points into different clusters by iteratively identifying points that reside on the borders of the cluster and removing them until separable areas of data remain. During the peeling process, the method creates associations between the peeled border points and points in the inner layers by estimating the local density of the data points. These associations are then used to link between the separable unpeeled data points and thus form the resulting clusters.

The main idea of the method is the peeling of the border points which ensures that the cores of near-by clusters are clearly separated before the data points are classified as separate clusters. We present a novel formulation for identifying border points, which we validated analytically on a simple setting of a uniformly-distributed cluster. Additionally, as we have shown above, our border peeling technique can be used to infer confidence values associated with the data points of each cluster. Unlike other methods, we do not make strong assumptions about the structure of the data points or their density distribution such as a single density peak or uniform density levels. Furthermore, the method has been shown to be stable in the sense that it is insensitive to the setting of hard coded parameters.

We have extensively analyzed the performance of the Border-Peeling method over large datasets for which the number of clusters is unknown, and there is no prior knowledge about its general structure or density distribution. As we have shown, Border-Peeling clustering outperforms state-of-the-art non-parametric clustering methods, even when their free parameters are fine-tuned to achieve their best performance. In particular, the performance of Border-Peeling in comparison to DBSCAN is interesting

since conceptually the two techniques have much in common, as both methods extract the core of the clusters that separate adjacent clusters, and then expand core points to the rest of the cluster. However, the difference in performance is intriguing. We attribute the better performance of Border-Peeling clustering to the fact that the core points are not defined globally, but through an iterative process that senses the local densities. The incremental peeling not only identifies the border points but also carefully associates them with points that seem to be closer to the core of the cluster. Furthermore, the locally adaptive approach is advantageous in sensing and avoiding over-segmentation of the clusters. Having said that, we believe that the more significant advantage of Border-Peeling over DBSCAN and other non-parametric clustering methods is that the method is insensitive to small variations in parameter values, and that those values can be easily set according to the characteristics of the data set, as can be seen in the experiments section.

All our results were generated with the same set of parameters, which is detailed in the paper. To control the maximum association distance between data points, we estimated λ (introduced in Section 2.2) according to the global behavior of the data. We believe that this global estimation can be further refined to better suit datasets with different structures and densities, perhaps in a semi-supervised setting where some samples are labeled. In general, as future work, we believe that our unsupervised technique can be expanded to accommodate supervised scenarios, to achieve improved accuracy in such controlled settings.

References

[1] R. J. Campello, D. Moulavi, and J. Sander. Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining*, pages 160–172. Springer, 2013. 5

[2] Y. Cheng. Mean shift, mode seeking, and clustering. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 17(8):790–799, 1995. 1

[3] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996. 1

[4] B. J. Frey and D. Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007. 5

[5] L. Fu and E. Medico. Flame, a novel fuzzy clustering method for the analysis of dna microarray data. *BMC bioinformatics*, 8(1):3, 2007. 2, 5

[6] A. Gionis, H. Mannila, and P. Tsaparas. Clustering aggregation. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):4, 2007. 5

[7] L. Hubert and P. Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985. 6

[8] I. Jolliffe. *Principal component analysis*. Wiley Online Library, 2002. 6

[9] F. Korn and S. Muthukrishnan. Influence sets based on reverse nearest neighbor queries. In *ACM SIGMOD Record*, volume 29, pages 201–212. ACM, 2000. 2

[10] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. 6

[11] Y. LeCun and C. Cortes. MNIST handwritten digit database. 2010. 6

[12] J. MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA., 1967. 1

[13] A. Y. Ng, M. I. Jordan, Y. Weiss, et al. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 2:849–856, 2002. 5

[14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. 5

[15] S. Van Der Walt, S. C. Colbert, and G. Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011. 5

[16] A. Vedaldi and K. Lenc. Matconvnet – convolutional neural networks for matlab. In *Proceeding of the ACM Int. Conf. on Multimedia*, 2015. 6

[17] C. J. Veenman, M. J. Reinders, and E. Backer. A maximum variance cluster algorithm. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(9):1273–1280, 2002. 5

[18] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *The Journal of Machine Learning Research*, 11:2837–2854, 2010. 6

[19] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007. 6

[20] C. Xia, W. Hsu, M. L. Lee, and B. C. Ooi. Border: efficient computation of boundary points. *Knowledge and Data Engineering, IEEE Transactions on*, 18(3):289–303, 2006. 2

[21] L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *Advances in neural information processing systems*, pages 1601–1608, 2004. 2